

The Public Option: a Non-regulatory Alternative to Network Neutrality

Richard T. B. Ma

Department of Computer Science
National University of Singapore
Email: tbma@comp.nus.edu.sg

Vishal Misra

Department of Computer Science
Columbia University, New York, USA
Email: misra@cs.columbia.edu

Abstract

Network neutrality and the role of regulation on the Internet have been heavily debated in recent times. Amongst the various definitions of network neutrality, we focus on the one which *prohibits paid prioritization of content* and we present an analytical treatment of the topic. We develop a model of the Internet ecosystem in terms of three primary players: consumers, ISPs and content providers. Our analysis looks at this issue from the point of view of the consumer, and we describe the desired state of the system as one which maximizes *consumer surplus*. By analyzing different scenarios of monopoly and competition, we obtain different conclusions on the desirability of regulation. We also introduce the notion of a *Public Option ISP*, an ISP that carries traffic in a network neutral manner. Our major findings are (i) in a monopolistic scenario, network neutral regulations benefit consumers; however, the introduction of a Public Option ISP is even better for consumers, as it aligns the interests of the monopolistic ISP with the consumer surplus and (ii) in an oligopolistic situation, the presence of a Public Option ISP is again preferable to network neutral regulations, although the presence of competing price-discriminating ISPs provides the most desirable situation for the consumers.

I. INTRODUCTION

Since around 2005, network neutrality has been a hotly debated topic amongst law and policy makers. The core debate has centered around the argument whether ISPs should be allowed to provide service differentiation and/or user discrimination, with the notion of “user” being either content providers (CPs) or consumers. Proponents of network neutrality, mostly the CPs, have argued that the Internet has been “neutral” since its inception and that has been a critical factor in the innovation and rapid growth that has happened on it. Opponents of network neutrality, mostly the ISPs, claim that without some sort of service differentiation, ISPs will lose the incentive to invest in the networks and the end user experience will suffer. Both camps implicitly or explicitly claim that their approach is beneficial for consumers. A recent Federal Communications Commission (FCC) vote [1] in the US has sided with the proponents, although the ruling leaves some room for service differentiation in wireless networks. The controversy rages on though with corporations like Verizon filing lawsuits challenging the ruling and a “toll-tax” dispute between Level3/Netflix and Comcast being cast as a network neutrality issue.

We study the issue explicitly from the point of view of the consumer under both monopolistic and oligopolistic scenarios. A lot of arguments against network neutrality live in an idealized world where economies of scale do not exist and monopolies cannot emerge, and therefore perfect competition solves all problems. We believe reality is more nuanced and hence we examine monopolistic scenarios as well.

Our approach to the analysis is a game theoretic one, and we focus on the consumer surplus. We model the rate allocation mechanism of the system and the user demand for different CPs. The interplay between the two determines the rate equilibrium for all traffic flows. Our model of price discrimination is for the ISPs to offer two classes of service to CPs. The ISP divides its capacity into a premium and ordinary class, and CPs get charged for carrying traffic in the premium class, and more details are presented in Section III-A. We then identify and analyze the strategic games played between ISPs, CPs and consumers in Section III for a monopolistic scenario and in Section IV for oligopolistic scenarios. In Section IV-A,

we introduce the notion of a *Public Option* ISP which is neutral to all CPs. The Public Option ISP can be implemented by processes like local loop unbundling [3] in a monopolistic market and either government or a private organization can run the ISP and still be profitable [12]. Given the framework, our major findings are:

- The impact of network neutrality on consumer surplus depends on the nature of competition at the ISP level. Concretely, a neutral network is beneficial for consumer surplus under a monopolistic regime (Section III), whereas a non-neutral network is advantageous for consumers under an oligopolistic scenario (Section IV).
- The capacity of the network plays a big role on the regulations for consumer surplus. Price discrimination in a monopolistic network where capacity is plentiful can have a damaging effect on consumer surplus, whereas it might gain higher utility for the consumer in a network where capacity is extremely scarce (Section III).
- Introducing a Public Option ISP is advantageous for consumers. In a monopolistic situation, the Public Option ISP offers the best scenario for consumers (Theorem 5), followed by network neutral regulations, and an unregulated market being the worst.
- In an oligopolistic situation, the Public Option ISP is still preferable to network neutral regulations; however, since the incentive for an ISP to gain market share is aligned with the maximization of consumer surplus (Theorem 6), no regulation is needed to protect the consumers.
- Under an oligopolistic competition, any ISP's optimal pricing and service differentiation strategy, whether network neutral or not, will be close the one that maximizes consumer surplus (Theorem 6 and Corollary 1). Moreover, under a probable equilibrium where ISPs use homogenous strategies, their market shares will be proportional to their capacities (Lemma 4), which implies that ISPs do have incentives to invest and expand capacity so as to their increase market shares.

Our paper sheds new light on the network neutrality debate and concretely identifies where and how regulation can help. Additionally, our identification of the Public Option ISP is especially important as it provides a solution that combines the best of both worlds, protecting consumer interests without enforcing strict regulations on all ISPs. We start with describing our model in the next section.

II. THREE-PARTY ECOSYSTEM MODEL

We consider a model of the Internet with three parties: 1) CPs, 2) ISPs and 3) consumers. We focus on a fixed consumer group in a targeted geographic region. We denote M as the number of consumers in the region¹. Each consumer subscribes to an Internet access service via an ISP. We consider the scenarios where one monopolistic ISP I or a set \mathcal{I} of competing oligopolistic ISPs provide the Internet access for the consumers. We denote \mathcal{N} as the set of CPs from which the consumers request content. We define $N = |\mathcal{N}|$ as the number of CPs. Our model does not include the backbone ISPs for two reasons. First, the bottleneck of the Internet is often at the last-mile connection towards the consumers [10], both wired and wireless. We focus on the regional or so-called *eyeball ISPs* [14] that provide the bottleneck last-mile towards the consumers. Second, the recent concern on network neutrality manifests itself in the cases where the last-mile ISPs, e.g. France Telecom, Telecom Italia and Vodafone, intended to differentiate services and charge CPs, e.g. Apple and Google, for service fees [5].

We denote μ as the last-mile bottleneck capacity towards the consumers in the region. Figure 1 depicts the contention at the bottleneck among different flows from the CPs. We denote λ_i as the aggregate throughput rate from CP i to the consumers. Because consumers initiate downloads and retrieve content from the CPs, we first model the consumer demands so as to characterize the CPs' throughput rates λ_i s.

¹Note that M can also be interpreted as the *average* or *peak* number of consumers accessing the Internet simultaneously in the region, which will scale with the total number of actual consumers in a region. This does not change the nature of any of the results we describe subsequently, but gives a more realistic interpretation of the rate equilibrium.




Fig. 1. Contention at the last-mile bottleneck link.

Given a set \mathcal{N} of CPs, a group of M consumers and a link with capacity μ , we denote the system as a triple (M, μ, \mathcal{N}) .

A. Throughput Demand

We denote $\hat{\theta}_i$ as the *unconstrained throughput* for a typical user of CP i . For instance, the unconstrained throughput for the highest quality Netflix streaming movie is about 5 Mbps [4], and given an average query page of 20 KB and an average query response time of .25 seconds [2], the unconstrained throughput for a Google search is about 600 Kbps, or just over $1/10^{th}$ of Netflix. We denote $\alpha_i \in (0, 1]$ as the percentage of consumers that ever access CP i 's content, which models the popularity of the content of CP i . We define $\hat{\lambda}_i = \alpha_i M \hat{\theta}_i$ as the unconstrained throughput of CP i . Without contention, CP i 's throughput λ_i equals $\hat{\lambda}_i$. However, when the capacity μ cannot support the unconstrained throughput from all CPs, i.e. $\mu < \sum_{i \in \mathcal{N}} \hat{\lambda}_i$, two things will happen: 1) a typical user of CP i obtains throughput $\theta_i < \hat{\theta}_i$ from CP i , and 2) some active users might stop downloading content from CP i when θ_i goes below certain threshold, e.g. users of real-time content like Netflix. We denote θ_i as the *achievable throughput* for the consumers downloading content from CP i . We define a demand function $d_i(\theta_i)$ which represents the percentage of consumers that still demand content from CP i under the achievable throughput θ_i .

Assumption 1: For any CP i , the demand $d_i(\cdot)$ is a non-negative, continuous and non-decreasing function defined on the domain of $[0, \hat{\theta}_i]$, and satisfies $d_i(\hat{\theta}_i) = 1$.

We define the aggregate throughput as $\lambda_{\mathcal{N}} = \sum_{i \in \mathcal{N}} \lambda_i$, where each CP i 's throughput λ_i is defined as follows:

$$\lambda_i(\theta_i) = \alpha_i M d_i(\theta_i) \theta_i. \quad (1)$$

B. Rate Allocation Mechanism

When multiple flows share the same bottleneck link, they compete for capacity. The rates allocated to the flows depend on the rate allocation mechanism being use in the system.

Definition 1: A *rate allocation mechanism* is a function that maps any fixed demand profile $\{d_i : i \in \mathcal{N}\}^2$ to an achievable throughput profile $\{\theta_i : i \in \mathcal{N}\}$.

A rate allocation mechanism can be a flow control mechanism, e.g. CBR and VBR mechanisms, under which the bottleneck link decides the rates for each flow in a centralized manner, or a window-based end-to-end congestion control mechanism, e.g. TCP, under which each flow maintains a sliding window and adapts the its size based on implicit feedback from the network, e.g. the round-trip delay. We consider generic rate allocation mechanisms and assume that the resulting rate allocation obeys the physical constraints of the system and satisfies some desirable properties.

Axiom 1: $\theta_i \leq \hat{\theta}_i$ for all $i \in \mathcal{N}$.

Axiom 2: $\lambda_{\mathcal{N}} = \min \left\{ \mu, \sum_{i \in \mathcal{N}} \hat{\lambda}_i \right\}$.

The above axioms characterize the feasibility of an allocation: the aggregate rate cannot exceed the capacity and the individual rate would not exceed its unconstrained throughput. It also characterizes a

²Without a bracket, we use d_i as a fixed demand.

work-conserving property: if congestion can be alleviated without increasing capacity μ , the allocation would do so by fully utilizing the capacity.

Axiom 3 (Monotonicity): A rate allocation is *monotonic* if for any $M > 0$ and capacity $\mu_1 < \mu_2$, the achievable throughput θ_i for any $i \in \mathcal{N}$ satisfies

$$\theta_i(M, \mu_2, \mathcal{N}) \geq \theta_i(M, \mu_1, \mathcal{N}).$$

The Monotonicity axiom implies that if a flow gets an achievable throughput in one system, it will be allocated at least that amount of throughput under a less congested system.

Axiom 4 (Independence of Scale): A rate allocation is *independent of scale* if for any $\xi > 0$, the achievable throughput θ_i for any $i \in \mathcal{N}$ satisfies

$$\theta_i(M, \mu, \mathcal{N}) = \theta_i(\xi M, \xi \mu, \mathcal{N}).$$

The Independence of Scale axiom states that if the capacity scales at the same rate as the consumer size, each flow's achievable throughput θ_i remains the same.

Assumption 2: The network system implements a rate allocation mechanism that satisfies Axiom 1 to 4.

C. Rate Equilibrium

The demand functions map the achievable throughput to a level of demand; the rate allocation mechanisms map fixed demands to achievable throughput. The interplay between a rate allocation mechanism and the demand functions determines the system rate equilibrium as the following theorem.

Theorem 1: A system (M, μ, \mathcal{N}) has a unique rate equilibrium $\{\theta_i : i \in \mathcal{N}\}$ under Assumption 1 and Axiom 1 to 3.

Proof of Theorem 1: Based on Assumption 1, we know that for any $i \in \mathcal{N}$, the throughput $\lambda_i(\theta_i)$ is a non-decreasing and continuous function. By Axiom 1, $\lambda_i(\theta_i)$ has a range of $[0, \hat{\lambda}_i]$. We show the uniqueness of the rate equilibrium by the following two cases.

We first consider the case where $\mu \geq \sum_{i \in \mathcal{N}} \hat{\lambda}_i$. By Axiom 2, $\lambda_{\mathcal{N}} = \sum_{i \in \mathcal{N}} \lambda_i(\theta_i) = \sum_{i \in \mathcal{N}} \hat{\lambda}_i$. Because $\lambda_i \leq \hat{\lambda}_i$ for all $i \in \mathcal{N}$, we must have $\lambda_i = \hat{\lambda}_i$. Therefore, the unique rate equilibrium must be $\{\theta_i = \hat{\theta}_i : i \in \mathcal{N}\}$.

We then consider the case where $\mu < \sum_{i \in \mathcal{N}} \hat{\lambda}_i$. By Axiom 2, $\lambda_{\mathcal{N}} = \sum_{i \in \mathcal{N}} \lambda_i(\theta_i) = \mu$. Because each $\lambda_i(\theta_i)$ is non-decreasing in θ_i and continuous in the range of $[0, \hat{\lambda}_i]$, we can always find a solution $\{\theta_i : i \in \mathcal{N}\}$ that satisfies $\sum_{i \in \mathcal{N}} \lambda_i(\theta_i) = \mu$. Suppose there exists two equilibrium solutions $\{\theta_i^1 : i \in \mathcal{N}\}$ and $\{\theta_i^2 : i \in \mathcal{N}\}$. For any $\varepsilon > 0$, we denote $\{\theta_i^3 : i \in \mathcal{N}\}$ as a rate equilibrium of the system $(M, \mu - \varepsilon, \mathcal{N})$. By Axiom 2, $\sum_{i \in \mathcal{N}} \lambda_i(\theta_i^3) = \mu - \varepsilon$. By Axiom 3, $\theta_i^3 \leq \min\{\theta_i^1, \theta_i^2\}$ for all $i \in \mathcal{N}$. Thus, we have

$$\begin{aligned} \mu - \varepsilon &= \sum_{i \in \mathcal{N}} \lambda_i(\theta_i^3) \leq \sum_{i \in \mathcal{N}} \lambda_i(\min\{\theta_i^1, \theta_i^2\}) \\ &= \sum_{i \in \mathcal{N}} \min\{\lambda_i(\theta_i^1), \lambda_i(\theta_i^2)\} < \mu. \end{aligned}$$

However, the above inequality would be violated when ε approaches zero. By showing this contradiction, we prove the uniqueness of the rate equilibrium under $\mu < \sum_{i \in \mathcal{N}} \hat{\lambda}_i$. ■

We define $\nu = \mu/M$ as the per capita capacity of the system. By Axiom 4, we further characterize the rate equilibrium $\{\theta_i : i \in \mathcal{N}\}$ as follows.

Lemma 1: Under Assumption 1 and 2, for all $i \in \mathcal{N}$, θ_i in equilibrium can be expressed as $\theta_i(M, \mu, \mathcal{N}) = \theta_i(\nu, \mathcal{N})$, which is a non-decreasing and continuous function in ν .

Proof of Lemma 1: By Axiom 4, $\theta_i(M_1, \mu_1, \mathcal{N})$ equals $\theta_i(M_2, \mu_2, \mathcal{N})$ if $M_1/\mu_1 = M_2/\mu_2$. Therefore, we can denote $\theta_i(\nu, \mathcal{N}) = \theta_i(M, \mu, \mathcal{N})$ for all systems (M, μ, \mathcal{N}) with $\nu = M/\mu$. Given any two systems

$(M_1, \mu_1, \mathcal{N})$ and $(M_2, \mu_2, \mathcal{N})$ with $\nu_1 \leq \nu_2$, we first show that $\theta_i(\nu_1, \mathcal{N}) \leq \theta_i(\nu_2, \mathcal{N})$ for all $i \in \mathcal{N}$. Suppose there exists $i \in \mathcal{N}$ such that $\theta_i(\nu_1, \mathcal{N}) > \theta_i(\nu_2, \mathcal{N})$. By Axiom 4, we have

$$\begin{aligned} \theta_i(M_1, \mu_1, \mathcal{N}) &= \theta_i(\nu_1, \mathcal{N}) \\ &> \theta_i(\nu_2, \mathcal{N}) = \theta_i(M_2, \mu_2, \mathcal{N}) = \theta_i(M_1, \frac{M_1}{M_2}\mu_2, \mathcal{N}). \end{aligned}$$

Because $\frac{M_1}{M_2}\mu_2 > \mu_1$, the above inequality violates Axiom 3. By showing this contradiction, we show that $\theta_i(\nu, \mathcal{N})$ is a non-decreasing function of ν . As a result, if $\theta_i(\nu, \mathcal{N})$ is not continuous, then it can only have upward jumps. Therefore, $\lambda_{\mathcal{N}} = \sum_{i \in \mathcal{N}} \lambda_i(\theta_i)$ will only have upward jumps as well. However, by Axiom 2, $\lambda_{\mathcal{N}}$ has to meet the equality constraint $\lambda_{\mathcal{N}} = \min \{ \mu, \sum_{i \in \mathcal{N}} \hat{\lambda}_i \}$, which implies that when ν increases from zero, $\lambda_{\mathcal{N}}$ cannot have jumps. Therefore, each θ_i has to be continuous in ν . ■

Lemma 1 states that when ν increases, users' achievable throughput θ_i would not worse off for any CP i .

We denote $\phi_i \geq 0$ as the per unit traffic utility that the consumers obtain by receiving content from CP i . This utility can be derived from communicating with friends, e.g. Skype, watching movies, e.g. Netflix, obtaining information, e.g. Google, RapidShare, and etc. Our model does not assume any form of this utility. We denote CS as the consumer surplus defined as $CS = \sum_{i \in \mathcal{N}} \phi_i \lambda_i$ and denote Φ as the per capita consumer surplus defined as

$$\Phi = \frac{CS}{M} = \frac{1}{M} \sum_{i \in \mathcal{N}} \phi_i \lambda_i(\theta_i) = \sum_{i \in \mathcal{N}} \phi_i \alpha_i d_i(\theta_i) \theta_i. \quad (2)$$

Theorem 2: Under Assumption 1 and 2, the per capita consumer surplus Φ can be expressed as $\Phi(M, \mu, \mathcal{N}) = \Phi(\nu, \mathcal{N})$, which is non-decreasing function in ν . In particular, it strictly increases in $\nu \in [0, \sum_{i \in \mathcal{N}} \alpha_i \hat{\theta}_i]$.

Theorem 2 states that the per capita consumer surplus will strictly increase with the system per capita capacity ν , unless it is already maximized when unconstrained throughput is obtained. This result does not depend on the values of ϕ_i s.

Proof of Theorem 2: If we define $\Phi_i = \phi_i \alpha_i d_i(\theta_i) \theta_i$ for all $i \in \mathcal{N}$, by definition $\Phi = \sum_{i \in \mathcal{N}} \Phi_i$. Since θ_i is a function of ν by Lemma 1, Φ_i and Φ are functions of ν as well. By Assumption 1, Φ_i is a non-decreasing function of θ_i . By Lemma 1, θ_i is non-decreasing in ν ; therefore, Φ_i and then Φ are non-decreasing in ν . By Axiom 2, $\lambda_{\mathcal{N}} = \mu$ when $\nu \in [0, \sum_{i \in \mathcal{N}} \alpha_i \hat{\theta}_i]$, which implies that when ν increases, there must exist some $i \in \mathcal{N}$ with θ_i strictly increasing. As a result, Φ_i and Φ have to be strictly increasing as well. ■

D. Examples and Illustrations

In this subsection, we illustrate some examples of demand functions and rate allocation mechanisms.

1) *Demand as a function of throughput sensitivity:* Distinct CPs often have different demand patterns. For example, the demand for real-time applications decreases dramatically when its throughput drops below certain threshold where performance cannot be tolerated by most of users. We can characterize this throughput sensitivity of the CPs by a positive parameter β_i and consider the demand function

$$d_i(\theta_i) = e^{-\beta_i \left(\frac{\hat{\theta}_i}{\theta_i} - 1 \right)} = e^{-\beta_i \left(\frac{1}{\omega_i} - 1 \right)}, \quad (3)$$

where we define $\omega_i = \theta_i / \hat{\theta}_i$ as the percentage of unconstrained throughput achieved for CP i . The user demand decays exponentially with the level of congestion (measured by $\frac{\hat{\theta}_i - \theta_i}{\theta_i}$, the ratio of unsatisfied and achieved throughput) scaled by β_i . This demand function distinguishes the CPs via their throughput sensitivity β_i : larger β_i indicates higher sensitivity to throughput for CP i . Figure 2 illustrates the demand functions with various values of β_i . To normalize $\hat{\theta}_i$, we plot d_i against ω_i instead of θ_i . We observe that




Fig. 2. Demand function $d_i(\omega_i)$.

when throughput drops linearly, the demand drops sharply for large β_i , e.g. when $\beta_i = 5$, the demand is halved with a 10% drop in throughput from $\hat{\theta}_i$. Large β_i s can be used to model CPs that have stringent throughput requirements, e.g. Netflix; while, small β_i s can be used to model CPs that are less sensitive to throughput, e.g. a Google search query.

2) *End-to-end congestion control mechanisms*: Due to the end-to-end design principle of the Internet [9], congestion control has been implemented by window-based protocols, i.e. TCP and its variations. Mo and Walrand [19] showed that a class of α -proportional fair solutions³ can be implemented by window-based end-to-end protocols. Among the class of α -proportional fair solutions, the max-min fair allocation, a special case with $\alpha = \infty$, is the result of the AIMD mechanism of TCP [7]. Differing round trip times, receiver window sizes and loss rates can result in different bandwidths, but to a first approximation, TCP provides a max-min fair allocation of available bandwidth amongst flows. Although other protocols, e.g. UDP, coexist in the Internet, recent research [16] sees a growing concentration of application traffic, especially video, over TCP.

We illustrate the rate allocation under the max-min fair mechanism using an example of three CPs with demand functions of Equation (3) and parameters $(\alpha_1, \hat{\theta}_1, \beta_1) = (1, 1, 0.1)$, $(\alpha_2, \hat{\theta}_2, \beta_2) = (0.3, 10, 3)$ and $(\alpha_3, \hat{\theta}_3, \beta_3) = (0.5, 3, 5)$. CP 1 represents Google-type of CPs that are extensively accessed and less sensitive to throughput. CP 2 represents Netflix-type of CPs that are more throughput-sensitive and have high unconstrained throughput. CP 3 represents Skype-type of CPs that are extremely sensitive to throughput and have medium unconstrained throughput.

Figure 3 illustrates the rates and the corresponding demands of the three CPs under a max-min fair allocation mechanism. We vary the per capita capacity ν from 0 to 6,000. We observe that when ν increases from zero, the demand for Google-type content increases first, followed by the demand for Skype-type content and, the demand for Netflix-type content being the last.

III. MONOPOLISTIC ISP ANALYSIS

In this section, we start with the scenario where the last-mile capacity is controlled by a single monopolistic ISP I . We analyze the ISP's strategy under which non-neutral service differentiation is allowed, and the corresponding best responses of the CPs. We derive the equilibria of the system and analyze the ISP's impact on the system congestion and the welfare of the consumers and the ISP itself.

³Any α -proportional fair solution also satisfies Assumption 2.




Fig. 3. Throughput under max-min fair mechanism.

A. Non-Neutral Service Differentiation

We assume that the monopolistic last-mile ISP I has a capacity of μ . This ISP can be a retail residential ISP, e.g. Comcast and Time Warner Cable, or a mobile operator, e.g. Verizon and AT&T. Regardless of whether it is a wired or wireless provider, it serves as the last-mile service provider for the consumers. We assume that the ISP is allowed to allocate a fraction $\kappa \in [0, 1]$ of its capacity to serve premium CPs and charge them at a rate $c \in [0, \infty)$ (dollar per unit traffic). For a wired ISP, κ can be interpreted as the percentage of capacity deployed for private peering points that charge a fee of c per unit incoming traffic and $1 - \kappa$ can be interpreted as the percentage of capacity deployed for public peering points where incoming traffic is charge-free. For a wireless ISP, κ can be interpreted as the percentage of capacity devoted for the premium traffic that will be charged at a rate of c . The pair of parameters (κ, c) can also be thought of a type of Paris Metro Pricing (PMP) [21], [23], where one ordinary and another premium service class have capacities of $(1 - \kappa)\mu$ and $\kappa\mu$ and charge 0 and c respectively. In reality, content might be delegated via content distribution networks (CDNs), e.g. Akamai, or backbone ISPs, e.g. Level3 is a major tier-1 ISP that delivers Netflix traffic towards regional ISPs. Therefore, in practice, the charge c might be imposed on the delivering ISP, e.g. Level3, and then be recouped from the CP, e.g. Netflix, by its delivering ISP, e.g. Level3. Our model does not assume any form of the implementation.

We denote \mathcal{O} and \mathcal{P} as the two disjoint sets of CPs that join the ordinary and premium class respectively. We denote v_i as CP i 's per unit traffic revenue. This revenue can be generated by advertising for media clients, e.g. Google, or by selling products to online consumers, e.g. Amazon, or by providing services to consumers, e.g. Netflix and e-banking. Our model does not assume how the revenue is generated either. Each CP i 's utility function u_i can be expressed as

$$u_i(\lambda_i) = \begin{cases} v_i \lambda_i & \text{if } i \in \mathcal{O}, \\ v_i \lambda_i - c \lambda_i & \text{if } i \in \mathcal{P}. \end{cases} \quad (4)$$

We define $IS = c \lambda_{\mathcal{P}}$ as the ISP surplus (CP-side revenue) and denote Ψ as the per capita ISP surplus defined as

$$\Psi = \frac{IS}{M} = \frac{c}{M} \lambda_{\mathcal{P}} = \frac{c}{M} \sum_{i \in \mathcal{P}} \lambda_i(\theta_i) = c \sum_{i \in \mathcal{P}} \alpha_i d_i(\theta_i) \theta_i.$$

B. Content Provider's Best Response

Given the ISP's decision κ and c , each CP chooses whether to join the ordinary service class \mathcal{O} or the premium class \mathcal{P} . We denote ρ_i as the per capita throughput over CP i 's user base, i.e. $\alpha_i M$ users,

defined as

$$\rho_i(\nu, \mathcal{N}) = d_i(\theta_i(\nu, \mathcal{N}))\theta_i(\nu, \mathcal{N}). \quad (5)$$

Lemma 2: Given a fixed set \mathcal{O} of CPs in the ordinary class and a fixed set \mathcal{P} of CPs in the premium class, a new CP i 's optimal strategy is to join the premium service class, if

$$(v_i - c) \rho_i(\kappa\nu, \mathcal{P} \cup \{i\}) \geq v_i \rho_i((1 - \kappa)\nu, \mathcal{O} \cup \{i\}). \quad (6)$$

Moreover, when equality reaches, CP i obtains the same utility in both service classes.

Proof of Lemma 2: When joining the ordinary service class, the throughput of CP i is

$$\begin{aligned} \lambda_i(M, (1 - \kappa)\mu, \mathcal{O} \cup \{i\}) &= \alpha_i M \rho_i(M, (1 - \kappa)\mu, \mathcal{O} \cup \{i\}) \\ &= \alpha_i M \rho_i((1 - \kappa)\nu, \mathcal{O} \cup \{i\}). \end{aligned}$$

When joining the premium service class, the throughput is

$$\begin{aligned} \lambda_i(M, \kappa\mu, \mathcal{P} \cup \{i\}) &= \alpha_i M \rho_i(M, \kappa\mu, \mathcal{P} \cup \{i\}) \\ &= \alpha_i M \rho_i(\kappa\nu, \mathcal{P} \cup \{i\}). \end{aligned}$$

Therefore, by Equation 4, CP i 's utility is

$$u_i = \begin{cases} v_i \alpha_i M \rho_i((1 - \kappa)\nu, \mathcal{O} \cup \{i\}) & \text{if } i \in \mathcal{O}, \\ (v_i - c) \alpha_i M \rho_i(\kappa\nu, \mathcal{P} \cup \{i\}) & \text{if } i \in \mathcal{P}. \end{cases}$$

By comparing the utilities that can be obtained in the two service classes, we obtain the condition (6). ■

Lemma 2 states that a CP will join the premium service class if that results higher profit, which is per-unit flow profit ($v_i - c$ for the premium class) multiplied by the per capita throughput ρ_i . The above decision is clear for a CP only if all other CPs have already made their choices. To treat all CPs equally, we model the decisions of all CPs as a simultaneous-move game as part of a two-stage game.

C. Two-Stage Strategic Game

We model the strategic behavior of the ISP and the CPs as a two-stage game, denoted as a quadruple (M, μ, \mathcal{N}, I) .

- 1) *Players:* The ISP I and the set of CPs \mathcal{N} .
- 2) *Strategies:* ISP I chooses a strategy $s_I = (\kappa, c)$. Each CP i chooses a binary strategy of whether to join the premium class. The CPs' strategy profile can be written as $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$, where $\mathcal{O} \cup \mathcal{P} = \mathcal{N}$ and $\mathcal{O} \cap \mathcal{P} = \emptyset$.
- 3) *Rules:* In the first stage, ISP I decides $s_I = (\kappa, c)$ and announces it to all the CPs. In the second stage, all the CPs make their binary decisions simultaneously and reach a joint decision $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$.
- 4) *Outcome:* The set \mathcal{P} of the CPs shares a capacity of $\kappa\mu$ and the set \mathcal{O} of the CPs shares a capacity of $(1 - \kappa)\mu$. Each CP $i \in \mathcal{O}$ gets a rate $\lambda_i(M, (1 - \kappa)\mu, \mathcal{O})$ and each CP $j \in \mathcal{P}$ gets a rate $\lambda_j(M, \kappa\mu, \mathcal{P})$.
- 5) *Payoffs:* Each CP i 's payoff is defined by the utility $u_i(\lambda_i)$ in Equation (4). The ISP's payoff is the revenue $IS = c\lambda_{\mathcal{P}}$ received from the premium class.

If we regard the set of CPs as a single player that chooses a strategy $s_{\mathcal{N}}$, our two-stage game can be thought of a Stackelberg game [22]. In this game, the first-mover ISP can take all the best-responses of the CPs into consideration and derive its optimal strategy s_I using *backward induction* [18]. Given any fixed strategy $s_I = (\kappa, c)$, the CPs derive their best strategies under a simultaneous-move game, denoted as $(M, \mu, \mathcal{N}, s_I)$. We denote $s_{\mathcal{N}}(M, \mu, \mathcal{N}, s_I) = (\mathcal{O}, \mathcal{P})$ as a strategy profile of the CPs under the game $(M, \mu, \mathcal{N}, s_I)$. Technically speaking, when $\kappa = 0$ or 1, there is only one service class. When $\kappa = 0$, we

define the trivial strategy profile as $s_{\mathcal{N}} = (\mathcal{N}, \emptyset)$; when $\kappa = 1$, although there is not a physical ordinary class, we define the trivial strategy profile as $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{N} \setminus \mathcal{O})$, with $\mathcal{O} = \{i : v_i \leq c, i \in \mathcal{N}\}$ which defines the set of ISPs that cannot afford to join the premium class. Based on Lemma 2, we can define an equilibrium in the sense of a Nash or competitive equilibrium. To break a tie, we assume that a CP always chooses to join the ordinary service class when both classes provide the same utility.

Definition 2: A strategy profile $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is a Nash equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, if

$$\frac{v_i - c}{v_i} \begin{cases} \leq \frac{\rho_i((1 - \kappa)\nu, \mathcal{O})}{\rho_i(\kappa\nu, \mathcal{P} \cup \{i\})} & \text{if } i \in \mathcal{O}, \\ > \frac{\rho_i((1 - \kappa)\nu, \mathcal{O} \cup \{i\})}{\rho_i(\kappa\nu, \mathcal{P})} & \text{if } i \in \mathcal{P}. \end{cases} \quad (7)$$

D. Competitive Equilibrium

Notice that a CP's joining decision to a service class might increase the congestion level and reduce the throughput of flows of that service class; however, if the number of CPs in a service class is big, an additional CP i 's effect will be marginal. Analogous to the *pricing-taking assumption* [18] in a competitive market, we can make a *throughput-taking assumption* for the CPs as follows.

Assumption 3: Any CP $i \notin \mathcal{N}$ makes an estimate $\tilde{\rho}_i(\nu, \mathcal{N})$ on its ex-post per capita throughput $\rho_i(\nu, \mathcal{N} \cup \{i\})$ in the decision-making under a competitive equilibrium.

Based on the above throughput-taking assumption, we can define a competitive equilibrium of the CPs as follows.

Definition 3: A strategy profile $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is a competitive equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, if

$$\frac{v_i - c}{v_i} \begin{cases} \leq \frac{\rho_i((1 - \kappa)\nu, \mathcal{O})}{\tilde{\rho}_i(\kappa\nu, \mathcal{P})} & \text{if } i \in \mathcal{O}, \\ > \frac{\tilde{\rho}_i((1 - \kappa)\nu, \mathcal{O})}{\rho_i(\kappa\nu, \mathcal{P})} & \text{if } i \in \mathcal{P}. \end{cases} \quad (8)$$

The competitive equilibrium depends on how each CP i calculates $\tilde{\rho}_i = d_i(\tilde{\theta}_i)\tilde{\theta}_i$, which boils down to an estimation of the ex-post throughput $\tilde{\theta}_i$. This estimation depends on the rate allocation mechanism being used. For example, under the max-min fair mechanism, CP i can expect an achievable throughput of $\theta_{\mathcal{N}} = \max\{\theta_j : j \in \mathcal{N}\}$. Thus, CP i can take this *throughput* as given and estimate that $\tilde{\theta}_i = \min\{\hat{\theta}_i, \theta_{\mathcal{N}}\}$.

To perform numerical evaluations, we will focus on competitive equilibria for two reasons. First, because the number of CPs in practice is big, the congestion-taking assumption is valid. Second, the *common knowledge* assumption [18] for reaching Nash equilibria might be questionable, because CPs rarely know the characteristics of all other CPs in practice. Nevertheless, most of our results apply for both equilibrium definitions. In the rest of the paper, unless we specifically indicate an equilibrium to be Nash or competitive, we use the term *equilibrium* to indicate both the Nash (Definition 2) and the competitive equilibrium (Definition 3).

Theorem 3: If $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is an equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, it is also a same type of equilibrium (Nash or competitive) of a game $(\xi M, \xi\mu, \mathcal{N}, s_I)$ for any $\xi > 0$.

Proof of Theorem 3: Under the same strategy $s_I = (\kappa, c)$, the new ordinary class $(\xi M, \xi(1 - \kappa)\mu, \mathcal{O})$ and the new premium class $(\xi M, \xi\kappa\mu, \mathcal{P})$ have the same per capita capacity $(1 - \kappa)\nu$ and $\kappa\nu$ respectively as before. By Lemma 1, the new system induces the same throughput θ_i as before. As a result, the solution

$(\mathcal{O}, \mathcal{P})$ will induce the same values of ρ_i and $\tilde{\rho}_i$, which is an estimate on ρ_i based on θ_i . Therefore, both sides of (7) and (8) do not change and the equilibrium conditions still hold. ■

Although a game $(M, \mu, \mathcal{N}, s_I)$ might have multiple equilibria, we do not assume that it reaches a particular equilibrium. However, to make equilibria under the same per capita capacity ν consistent, we make the following assumption.

Assumption 4: If $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ is a realized equilibrium of a game $(M, \mu, \mathcal{N}, s_I)$, then it is also the realized equilibrium of the linearly scaled game $(\xi M, \xi \mu, \mathcal{N}, s_I)$ for any $\xi > 0$.

The above assumption implies that when the ISP scales its capacity μ linearly and smoothly with its consumer size M , the CPs will not diverge abruptly into another equilibrium, if there exists any. For the game $(M, \mu, \mathcal{N}, s_I)$ with strategy $s_I = (\kappa, c)$ and equilibrium $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$, the per capita consumer surplus Φ is a function of ν , defined as

$$\Phi(M, \mu, \mathcal{N}, s_I) = \Phi(\nu, \mathcal{N}, s_I) = \Phi((1 - \kappa)\nu, \mathcal{O}) + \Phi(\kappa\nu, \mathcal{P}).$$

Under the above assumption, the per capita consumer and ISP surplus will remain the same in linearly scaled games $\{(\xi M, \xi \mu, \mathcal{N}, s_I) : \xi > 0\}$ in equilibrium.

Lemma 3: Under Assumption 4, the per capita consumer surplus Φ satisfies

$$\Phi(\nu, \mathcal{N}, s_I) = \Phi(\xi M, \xi \mu, \mathcal{N}, s_I), \quad \forall \xi > 0.$$

The above is true for the per capita ISP surplus Ψ as well.

Proof of Lemma 3: Given $s_I = (\kappa, c)$ and $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$, the per capita consumer surplus Φ in system (M, μ, \mathcal{N}) is

$$\Phi(M, \mu, \mathcal{N}, s_I) = \Phi(M, (1 - \kappa)\mu, \mathcal{O}) + \Phi(M, \kappa\mu, \mathcal{P}).$$

By Theorem 2, we can rewrite the above as

$$\Phi(M, \mu, \mathcal{N}, s_I) = \Phi((1 - \kappa)\nu, \mathcal{O}) + \Phi(\kappa\nu, \mathcal{P}).$$

By Assumption 4, we know that a linearly scaled system $(\xi M, \xi \mu, \mathcal{N})$ under the same strategy s_I will induce the same equilibrium $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$. As a result, we have

$$\Phi(\xi M, \xi \mu, \mathcal{N}, s_I) = \Phi((1 - \kappa)\nu, \mathcal{O}) + \Phi(\kappa\nu, \mathcal{P}) \quad \forall \xi > 0.$$

Since the right hand side of the above equations are the same, we can express Φ as a function of ν as

$$\Phi(\nu, \mathcal{N}, s_I) = \Phi(\xi M, \xi \mu, \mathcal{N}, s_I) \quad \forall \xi > 0.$$

Since $\Psi = c \sum_{i \in \mathcal{P}} \alpha_i \rho_i(\nu, \mathcal{N})$, under a scaled system where ν and \mathcal{P} do not change, Ψ does not change either. ■

E. Monopolistic ISP's Strategy

In order to generate revenue, the ISP's optimal strategy would encourage more CPs to join its premium service class.

Theorem 4: In the game (M, μ, \mathcal{N}, I) , for any $0 \leq c < 1$, strategy $s_I = (\kappa, c)$ is always dominated by $s_I^1 = (1, c)$. If $\lambda_{\mathcal{P}} < \min\{\mu, \sum_{v_i \geq c} \hat{\lambda}_i\}$, s_I is strictly dominated by s_I^1 . $s_I = (\kappa, c)$ is also dominated by $s_I' = (\kappa', c)$ with $\kappa' > \kappa$, if equilibrium $(\mathcal{O}', \mathcal{P}')$ under s_I' satisfies $\mathcal{P} \subseteq \mathcal{P}'$.

Proof of Theorem 4: Under $s_I^1 = (1, c)$, only the premium service class is provided and CPs will join the premium service class only if $v_i \geq c$. Therefore, the aggregate rate $\lambda_{\mathcal{P}}^1 = \min\{\mu, \sum_{i \in \mathcal{P}_c} \hat{\lambda}_i\}$, where $\mathcal{P}_c = \{i : v_i \geq c\}$. When keeping the same c , $\mathcal{P} \subseteq \mathcal{P}_c$ under any strategy s_I . We have

$$\lambda_{\mathcal{P}} = \min\{\kappa\mu, \sum_{i \in \mathcal{P} \subseteq \mathcal{P}_c} \hat{\lambda}_i\} \leq \min\{\mu, \sum_{i \in \mathcal{P}_c} \hat{\lambda}_i\} = \lambda_{\mathcal{P}}^1.$$

The above implies that the revenue $c\lambda_{\mathcal{P}} \leq c\lambda_{\mathcal{P}}^1$. Therefore, s_I is dominated by s_I^1 . If $\lambda_{\mathcal{P}} < \min\{\mu, \sum_{i \in \mathcal{P}_c} \hat{\lambda}_i\}$, $\lambda_{\mathcal{P}} < \lambda_{\mathcal{P}}^1$, and therefore s_I is strictly dominated by s_I^1 .

Similarly, for any $\kappa' > \kappa$ with $\mathcal{P} \subseteq \mathcal{P}'$, we have

$$\lambda_{\mathcal{P}} = \min\{\kappa\mu, \sum_{i \in \mathcal{P} \subseteq \mathcal{P}'} \hat{\lambda}_i\} \leq \min\{\kappa'\mu, \sum_{i \in \mathcal{P}'} \hat{\lambda}_i\} = \lambda_{\mathcal{P}'}.$$

Therefore, s_I is also dominated by s_I^1 . ■

When the monopoly ISP increases κ , it improves the condition in the premium service class and in a new equilibrium, \mathcal{P}' would attract more CPs to join than \mathcal{P} . Theorem 4 states that the ISP would have incentives to increase κ so as to maximize revenue. The effect of increasing κ is twofold: 1) more capacity is allocated to the premium class for sale, and 2) the reduced capacity in the ordinary class makes more ISPs switch to the premium class. As a result, one of the optimal strategies of the monopolistic ISP is to always set $\kappa = 1$. This implies that, if allowed, the selfish ISP will only provide a charged service class \mathcal{P} without contributing any capacity for the ordinary class \mathcal{O} . Suppose the ISP is allowed to set $\kappa = 1$, we first study its optimal price c and its impact on the consumer and ISP surplus.

We use the demand function of Equation (3) and the max-min fair mechanism for our numerical simulations. We study a scenario of 1000 CPs, whose α_i , $\hat{\theta}_i$ and v_i are uniformly distributed within $[0, 1]$ and β_i is uniformly distributed within $[0, 10]$. To satisfy all unconstrained throughput for the CPs, the per capita capacity needs to be around $\nu = 250$. Since throughput-sensitive applications, e.g. Skype, bring more utility to consumers in reality, we consider the consumer utility ϕ_i that is uniformly distributed within $[0, \beta_i]$ (the uniform distribution biases utility towards CPs with high throughput sensitivity while introducing some randomness)⁴.




Fig. 4. Per capita surplus Ψ and Φ under $\kappa = 1$.

Figure 4 plots Ψ and Φ versus the ISP's pricing strategy c under various per capita capacity ν ranging from 20 to 200. We observe three pricing regimes.

- 1) When c is small, Ψ increases linearly, i.e. $\Psi = c\nu$. This happens when most of the CPs can afford to join the service and the entire capacity is fully utilized, i.e. $\lambda_{\mathcal{P}} = \mu$, resulting a high level of consumer surplus Φ .
- 2) When c is large, Ψ drops sharply. This happens when only a small set of CPs can afford to join the service and the capacity is largely under-utilized, i.e. $\lambda_{\mathcal{P}} < \mu$, resulting a sharp drop in Φ accordingly.
- 3) When ν is abundant, e.g. $\nu = 200$, there exists a pricing region where Ψ increases sub-linearly and Φ decreases. Consequently, the ISP's optimal strategy ($c \approx 0.45$) could intentionally keep more

⁴We also used another setting of $\phi_i = U[0, U[0, 10]]$, that has the same scale but is independent of β_i s for all the experiments in Section III and IV. All the results are similar with the current setting and are shown in the appendix.




Fig. 5. Per capita surplus Ψ and Φ under various strategies $s_I = (\kappa, c)$ versus per capita capacity ν .

CPs away from the (only) service class and under-utilizes the capacity, which hurts the consumer surplus Φ .

As a result, the consumer surplus could be misaligned with a monopoly ISP's revenue when the capacity is abundant.

Figure 5 illustrates Ψ and Φ under various strategies $s_I = (\kappa, c)$ versus ν ranging up to 500, which doubles the required capacity to serve all unconstrained throughput. For a fixed c in each column, we identify three equilibrium regimes.

- 1) When ν is small, Ψ increases linearly and Φ increases accordingly. This happens when the premium class capacity is fully utilized, i.e. $\lambda_P = \kappa\mu$.
- 2) When ν keeps increasing, Ψ starts to decrease and Φ increases at a much slower rate. This happens when the premium class capacity is not fully utilized, i.e. $\lambda_P < \kappa\mu$, and more CPs move from \mathcal{P} to \mathcal{O} .
- 3) When ν is large, Ψ drops to zero for small values of κ , where Φ is maximized. This happens when $\mathcal{P} = \emptyset$ and \mathcal{O} 's capacity is abundant enough to serve all CPs' unconstrained throughput. However, if κ is big, e.g. $\kappa = 0.9$, it guarantees some revenue for the ISP, but reduces the consumer surplus from its maximum.

Further, under a fixed ν , we observe that higher κ induces higher revenue for the ISP (Theorem 4), even if that results an under-utilization of the premium class capacity and hurts the consumer surplus. When comparing different prices c , we observe that larger values of c make the premium class becomes under-utilized faster, because fewer CPs can afford to join the premium class when necessary. However, when reaching the turning point where congestion starts to be relieved, κ plays a major role, under which Φ 's rate of increase depends on how much percentage $1 - \kappa$ of capacity is allocated to the ordinary class \mathcal{O} . Notice that only under the exceptional case where 1) the capacity is extremely scarce, and 2) CPs with high values of v_i s also have high values of ϕ_i s for consumer utility, the ISP's optimal strategy of $\kappa = 1$ and large value of c might benefit the consumer surplus.

Under a fixed strategy s_I , $\Phi(\nu, \mathcal{N}, s_I)$ is not strictly non-decreasing in ν compared to the result of Theorem 2. The reason is that when ν varies, CPs might move between the service classes. In general, when ν changes a small amount such that $s_{\mathcal{N}} = (\mathcal{O}, \mathcal{P})$ does not change, the monotonicity still holds; however, when ν keeps increasing, CPs will move from \mathcal{P} to \mathcal{O} , upon which Φ might drop at the spot. We characterize this discontinuity by the following metric.

$$\epsilon_{s_I} = \sup\{\Phi(\nu_1, \mathcal{N}, s_I) - \Phi(\nu_2, \mathcal{N}, s_I) : \nu_1 < \nu_2\}. \quad (9)$$

Notice that ϵ_{s_I} captures the largest vertical distance of a downward gap in the curve $\Phi(\nu, \mathcal{N}, s_I)$. From Figure 5, we observe that when $|\mathcal{N}|$ is large, ϵ_{s_I} is quite small, which indicates the general trend of increasing Φ with ν .

Regulatory Implications: In the monopolistic scenario, the ISP with fixed capacity μ will maximize the percentage of capacity for a charged service class (κ) even if it is under-utilized. This implies that the ISP has the incentive to degrade service quality or avoid network upgrades or investments for the non-charged service class. To remedy this problem, the network neutrality principle should be imposed to some extent to protect consumer surplus. In other words, the non-neutral service differentiation should be limited, especially when the capacity μ is not scarce, by two means: 1) limit the percentage of capacity devoted to a charged premium class⁵, i.e. κ cannot be too large, such that the CPs in the ordinary class can obtain an appropriate amount of capacity, and 2) limit the charge c so that enough CPs would be able to join the premium class. The bottom line is that capacity under-utilization should be avoided under any case.

IV. OLIGOPOLISTIC ISP ANALYSIS

In the previous section, we concentrate on a monopolistic ISP I that has a capacity μ and uses a strategy $s_I = (\kappa, c)$. In this section, we extend our model to a set \mathcal{I} of oligopolistic ISPs, each $I \in \mathcal{I}$ of which has a capacity μ_I and uses a strategy $s_I = (\kappa_I, c_I)$. We define $\mu = \sum_{I \in \mathcal{I}} \mu_I$ as the total system capacity, and $m_I = M_I/M$ and $\gamma_I = \mu_I/\mu$ as the market share and capacity share of ISP I . Our oligopolistic models have two major differences with the monopolistic model. First, since consumers connect to the Internet via one of the ISPs, they might make strategic decisions on which ISP to subscribe to. We denote M_I as the consumer size of each ISP $I \in \mathcal{I}$, where $\sum_{I \in \mathcal{I}} M_I$ equals the total consumer size M . Second, besides maximizing the premium service revenue from the CPs, a more important objective of any ISP $I \in \mathcal{I}$ is to maximize their market share m_I of the consumers. This is what the last-mile ISPs can leverage on to generate the CP-side of the revenue in the first place.

Similar to the monopolistic ISP game (M, μ, \mathcal{N}, I) , we denote $(M, \mu, \mathcal{N}, \mathcal{I})$ as the two-stage oligopolistic ISP game, under which the set of ISPs \mathcal{I} choose their strategies $s_{\mathcal{I}} = \{s_I : I \in \mathcal{I}\}$ simultaneously in the first stage, and then the set of CPs \mathcal{N} and the M consumers make their strategic decisions simultaneously in a second-stage game $(M, \mu, \mathcal{N}, s_{\mathcal{I}})$. In the second-stage game, we denote $s_M = \{M_I : I \in \mathcal{I}\}$ as the consumers' strategy that determines all ISPs' market shares, and $s_{\mathcal{N}} = \{s_{\mathcal{N}}^I = (\mathcal{O}_I, \mathcal{P}_I) : I \in \mathcal{I}\}$ as the CPs' strategy, where each $s_{\mathcal{N}}^I$ denotes the decision made at ISP I .

We denote Φ_I as the per capita consumer surplus achieved at ISP I , defined as $\Phi_I(M_I, \mu_I, \mathcal{N}, s_I) = \Phi_I(\nu_I, \mathcal{N}, s_I) = \Phi((1 - \kappa_I)\nu_I, \mathcal{O}_I) + \Phi(\kappa_I\nu_I, \mathcal{P}_I)$, where $\nu_I = \mu_I/M_I$. We assume that consumers will move towards the ISPs that provide higher per capita surplus as follows.

Assumption 5: Under any fixed strategy profile $s_{\mathcal{I}}$ and $s_{\mathcal{N}}$, for any pair of ISPs $I, J \in \mathcal{I}$, consumers will move from I to J if $\Phi_I < \Phi_J$. This process stops when $\Phi_I = \Phi_J \forall I \in \mathcal{I}$ for some system-wise per capita consumer surplus $\Phi_{\mathcal{I}}$.

Although consumers might not be totally elastic or/and accessible to all available ISPs in practice, our assumption takes a macro perspective and assumes that if an ISP provides worse user-experience

⁵This was also suggested as a regulatory tool by an independent work of Shetty et. al. [23].

on average, there must exist some consumers who can and will move to other better ISPs. Based on Assumption 5, we define the equilibrium of the second-stage game $(M, \mu, \mathcal{N}, s_{\mathcal{I}})$ as follows.

Definition 4: A strategy profile $(s_M, s_{\mathcal{N}})$ is an equilibrium of the multi-ISP game $(M, \mu, \mathcal{N}, s_{\mathcal{I}})$ if 1) for any $I \in \mathcal{I}$, $s_{\mathcal{N}}^I$ is an equilibrium of the single-ISP game $(M_I, \mu_I, \mathcal{N}, s_I)$, and 2) $\Phi_I = \Phi_J$ for any $I, J \in \mathcal{I}$.

A. Duopolistic ISP Game

We first study a two-ISP game with $\mathcal{I} = \{I, J\}$. Before that, we formally define a Public Option ISP as follows.

Definition 5: A Public Option ISP PO is an ISP that uses a fixed strategy $s_{PO} = (0, 0)$ and does not divide its capacity or charge the CPs.

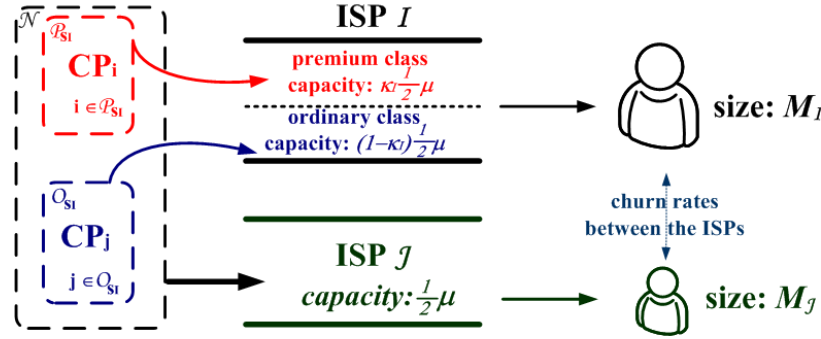


Fig. 6. A two-ISPs model.

We assume that ISP J is a Public Option ISP. Figure 6 illustrates an example of the above duopolistic model, where both ISPs have the same amount of capacity, the CPs choose a service class at ISP I and the consumers move between the ISPs. The above setting of the duopolistic game applies for two real scenarios. First, it models the competition between two ISPs, where one of them is actively a Public Option ISP and the other actively manages a non-neutral service differentiation. Second, it also models a situation where a single ISP owns the entire last-mile capacity μ ; however, by regulation [3], it is required to lease its capacity to other service providers that do not own the physical line. The leasing ISP might be technologically limited from providing service differentiation on the leased capacity, but actually have customers in the region. For both scenarios, we will answer 1) whether the non-neutral ISP could obtain substantial advantages over the neutral Public Option ISP (or whether the Public Option could survive under competition), and 2) how the competition is going to impact the consumer surplus.

We study the same set of 1000 CPs as in the previous section. We further assume that $\mu_I = \mu_J = \mu/2$ in our numerical example. We take the same route to numerically evaluate the competitive equilibria of the CPs under $\kappa_I = 1$.

Figure 7 plots ISP I 's market share m_I , per capita surplus Ψ_I , defined as $\Psi_I = \frac{c}{M} \lambda_{P_I}$, and Φ versus ISP I 's charge c_I . By the same reasons as before, the revenue of I increases linearly when its capacity is fully utilized, i.e. $\lambda_{P_I} = \kappa_I \mu_I$. However, we observe three differences: 1) after λ_{P_I} drops below $\kappa_I \mu_I$, Ψ_I drops to zero much steeper than before, 2) Φ_I does not drop down to zero when c_I increases to 1, and 3) the maximum Ψ_I is lower in the case of $\nu = 200$ than in the case of $\nu = 150$, which means that under $\kappa_I = 1$, capacity expansion could reduce ISP I 's revenue from the CPs. All these observations can be explained by checking at the market share of ISP I in the left sub-figure. The market share m_I starts to increase with c_I until ISP I 's capacity becomes under-utilized, i.e. $\lambda_{P_I} < \kappa_I \mu_I$. Afterwards, the market share drops dramatically. This explains that under congestion, i.e. $\lambda_{P_I} = \kappa_I \mu_I$, by increasing c_I , ISP I restricts the number of CPs in its service class and maintains less congestion, which could result higher




Fig. 7. ISP I 's market share m_I and per capita surplus Ψ_I and per capita consumer surplus Φ under $\kappa = 1$.

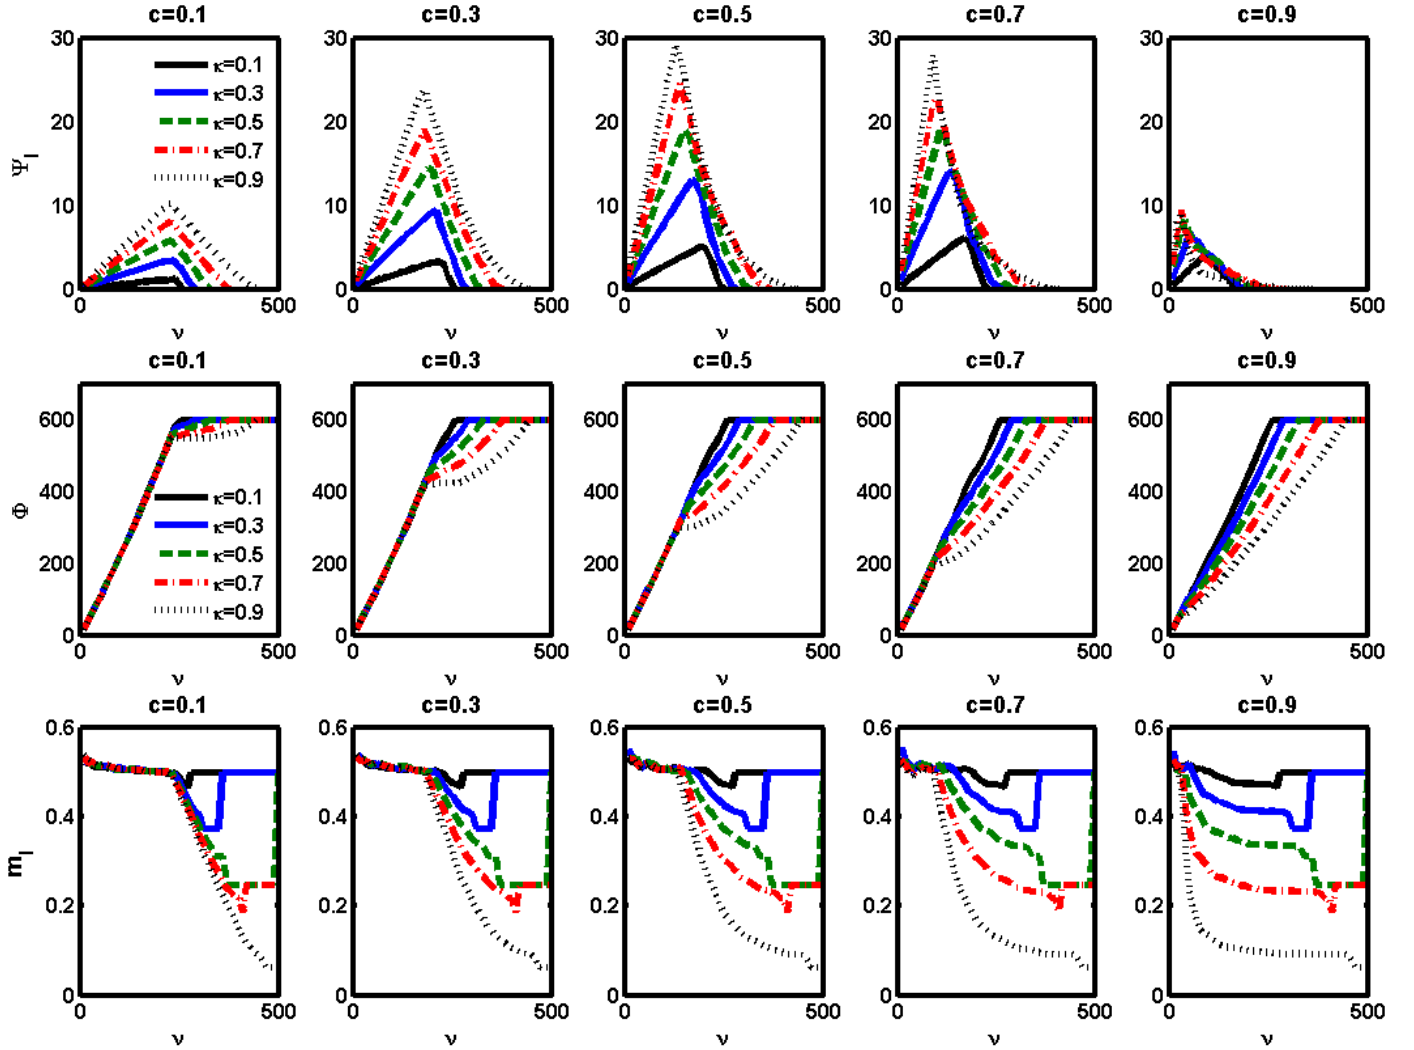


Fig. 8. Per capita surplus Ψ , Φ and market share m_I under various strategies $s_I = (\kappa, c)$ vs. per capita capacity ν .

consumer surplus, and therefore, attract more consumers from ISP J . After λ_{P_I} drops below $\kappa_I \mu_I$, further increase of c_I reduces the number of CPs in the service as well as the total throughput. This reduces consumer surplus, and therefore, consumers start to depart from ISP I to J . When c_I reaches 1, no CP survives in I 's service class and all consumers move to ISP J , which guarantees a non-zero consumer

surplus in equilibrium.

Parallel to Figure 5, Figure 8 illustrates the per capita surplus Ψ_I , Φ and ISP I 's market share m_I under various strategies s_I versus ν ranging up to 500. Compared to the monopolistic case, we observe two differences in Ψ_I and Φ : 1) under any strategy s_I , ISP I 's revenue drops sharply to zero after reaching a maximum point where λ_{P_I} drops below $\kappa_I \mu_I$, and 2) the increase of consumer surplus does not get affected by ISP I 's strategy too much. By observing the market share of ISP I , we identify two capacity regimes. First, when ν is extremely scarce, the differential pricing slightly benefits the consumer; and therefore, ISP I can obtain a slightly larger percentage of the market ⁶. Second, when the per capita capacity ν is abundant, ISP I obtains at most an equal share of the market if it uses a small value of κ . Under this case, the capacity under \mathcal{O} can support half of the population's unconstrained throughput and in fact, the premium class is empty, i.e. $\mathcal{P} = \emptyset$. As a result, ISP I follows the Public Option ISP by using some kind of neutral policy (small κ) and maximizes the consumer surplus.

Theorem 5: In the duopolistic game $(M, \mu, \mathcal{N}, \mathcal{I})$, where an ISP J is a Public Option, i.e. $s_J = (0, 0)$, if s_I maximizes M_I under an equilibrium (s_M, s_N) , then the per capita consumer surplus $\Phi_{\mathcal{I}}$ is also maximized under that equilibrium.

Proof of Theorem 5: For any strategy $s'_I \neq s_I$, we have two cases to analyze. First, if $M'_I = M_I$, then $M'_J = M - M'_I = M - M_I = M_J$. Given the same market share for the public option ISP, it induces the same per capita consumer surplus $\Phi'_J = \Phi_J$. In equilibrium, we have $\Phi'_{\mathcal{I}} = \Phi'_J = \Phi_J = \Phi_{\mathcal{I}}$. Second, if $M'_I < M_I$, then $M'_J = M - M'_I > M - M_I = M_J$. Given a larger market share for the public option ISP, the per capita capacity reduces, i.e. $\nu'_J < \nu_J$. By Theorem 2, the public option will not induce a larger per capita consumer surplus, i.e. $\Phi'_J \leq \Phi_J$. Thus, we have $\Phi'_{\mathcal{I}} = \Phi'_J \leq \Phi_J = \Phi_{\mathcal{I}}$ in equilibrium. ■

Theorem 5 implies that the existence of a Public Option ISP is superior to a network neutral situation, where $s_I = (0, 0)$. This is because given the freedom of choosing an optimal s_I to maximize market share, ISP I 's strategy will induce a maximum consumer surplus under $s_J = (0, 0)$.

Based on our results, we answer the previously raised two questions: 1) The non-neutral ISP cannot win substantially over the Public Option ISP, which can still be profitable under the competition, confirming the independent findings from [12]. 2) Regardless of the capacity size, the competition induces higher consumer surplus in equilibrium than under network neutral regulations. The strategic ISP could obtain slightly over 50% of the market; however, if it differentiates services in the way that hurts consumer surplus, its market share will drop sharply.

Regulatory Implications: In the duopolistic scenario with one of the ISPs being a Public Option, contrary to the monopolistic case, the non-neutral strategy s_I is always aligned with the consumer surplus (Theorem 5). This result shows an interesting alternative to remedy the network neutrality issue under a monopolistic market. Instead of enforcing the ISP to follow network neutrality, the government (or a private organization, since we independently verify it can be profitable) can provide the consumers with a Public Option ISP that is neutral to all CPs. Given such a neutral entity in the market, consumers will move to their public option if it provides higher consumer surplus than the non-neutral ISP that uses differential pricing to the CPs. Meanwhile, in order to maximize its market share, the non-neutral ISP will adapt its strategy to maximize consumer surplus. In conclusion, the introduction of a Public Option ISP is superior to network neutral regulations under a monopolistic market, since its existence aligns the non-neutral ISP's selfish interest with the consumer surplus.

B. Oligopolistic ISP Competition Game

After analyzing the duopolistic game between a non-neutral and a Public Option ISP, we further consider a deregulated market under which all ISPs make non-neutral strategies. We consider a multi-ISP game

⁶Notice that this is because by limiting the number of CPs in the premium class, the proportion of throughput-sensitive traffic is larger, which gives higher utility to consumers.

under which each ISP I chooses a strategy s_I to maximize its market share m_I .

We first consider a homogenous strategy $s = (\kappa, c)$, which can be a preferred or regulated strategy, used by all ISPs.

Lemma 4: If $s_{\mathcal{I}} = \{s_I = s : I \in \mathcal{I}\}$ for some strategy $s = (\kappa, c)$, then $\{m_I = \gamma_I, s_{\mathcal{N}}^I = s_{\mathcal{N}}(M, \mu, \mathcal{N}, s) : I \in \mathcal{I}\}$ is an equilibrium of the game (M, ν, \mathcal{N}, s) .

Lemma 4 shows a symmetric equilibrium where market share m_I is proportional to capacity μ_I . It implies that ISPs will have incentives to invest and expend capacity so as to obtain a larger market share. This equilibrium could be reached when ISPs simply mimic one another's strategy. A further question is whether the competition of market share among the ISPs would induce equilibria where consumer surplus is high. We denote s_{-I} as the strategy profile of the ISPs other than ISP I . Similar to the definition of ϵ_{s_I} in Equation (9), we define $\delta_{s_I} = \sup\{m_1 - m_2 : \Phi(\nu_1, \mathcal{N}, s_I) \leq \Phi(\nu_2, \mathcal{N}, s_I)\}$ and $\epsilon_{s_{-I}} = \max\{\epsilon_{s_J} : J \in \mathcal{I} \setminus \{I\}\}$.

Proof of Lemma 4: When $M_I = \gamma_I M$ and $s_I = s$ for all $I \in \mathcal{I}$, the single-ISP game $(M_I, \mu_I, \mathcal{N}, s)$ is a linearly scaled game of (M, μ, \mathcal{N}, s) . By Theorem 3, we know that

$$s_{\mathcal{N}}^I = s_{\mathcal{N}}(M, \mu, \mathcal{N}, s) = s_{\mathcal{N}}(M_I, \mu_I, \mathcal{N}, s_I), \quad \forall I \in \mathcal{I}.$$

By Lemma 3, we know that

$$\Phi_I = \Phi(M_I, \mu_I, \mathcal{N}, s_I) = \Phi(M, \mu, \mathcal{N}, s), \quad \forall I \in \mathcal{I}.$$

The above satisfies the two conditions of an equilibrium in Definition 4 and concludes the proof. \blacksquare

Theorem 6: Under any fixed strategy profile s_{-I} , if I 's strategy s_I is a best-response to s_{-I} that maximizes its market share m_I in the game $(M, \mu, \mathcal{N}, s_{\mathcal{I}})$, then s_I is a $\epsilon_{s_{-I}}$ -best-response for the per capita consumer surplus $\Phi_{\mathcal{I}}$, i.e.

$$\Phi_{\mathcal{I}} \geq \Phi'_{\mathcal{I}} - \epsilon_{s_{-I}}, \quad \forall s'_I \neq s_I.$$

Moreover, if s_I is a best-response that maximizes consumer surplus $\Phi_{\mathcal{I}}$ in the game $(M, \mu, \mathcal{N}, s_{\mathcal{I}})$, then s_I is a δ_{s_I} -best-response for the market share m_I , i.e.

$$m_I \geq m'_I - \delta_{s_I}, \quad \forall s'_I \neq s_I.$$

Proof of Theorem 6: If s_I maximizes the market share M_I , for any strategy $s'_I \neq s_I$, we have $M'_I \leq M_I$. Therefore, there exist an ISP $J \neq I$ such that $M'_J \geq M_J$. This implies $\nu'_J \leq \nu_J$. By the definition of ϵ_{s_J} in Equation (9), we have $\Phi'_J - \Phi_J \leq \epsilon_{s_J}$, or equivalently $\Phi_J \geq \Phi'_J - \epsilon_{s_J}$. Thus,

$$\Phi_{\mathcal{I}} = \Phi_J \geq \Phi'_J - \epsilon_{s_J} = \Phi'_{\mathcal{I}} - \epsilon_{s_J} \geq \Phi'_{\mathcal{I}} - \epsilon_{s_{-I}}.$$

If s_I maximizes the consumer surplus $\Phi_{\mathcal{I}}$, for any strategy $s'_I \neq s_I$, we have $\Phi'_{\mathcal{I}} \leq \Phi_{\mathcal{I}}$. By the definition of δ_{s_I} , we have $m'_I - m_I \leq \delta_{s_I}$, or equivalently $m_I \geq m'_I - \delta_{s_I}$. \blacksquare

Theorem 6 states that, given the fixed strategies of all other ISPs, an ISP's best-responses to maximize its market share and to maximize the consumer surplus are closely aligned. Parallel to Theorem 5, it shows that an ISP's selfish interest is, although not perfectly, aligned with the consumer surplus under competition. Technically, the $\epsilon_{s_{-I}}$ imperfection is due to the discontinuity of $\Phi(\nu, \mathcal{N}, s_I)$ in ν . When $\epsilon_{s_{-I}}$ approaches zero, Φ will be non-decreasing mostly and the objectives of maximizing market share and maximizing consumer surplus will converge.

Definition 6: A strategy profile $s_{\mathcal{I}} = \{s_I : I \in \mathcal{I}\}$ is a *market share* Nash equilibrium of the game $(M, \mu, s_{\mathcal{N}}, \mathcal{I})$ if for any $I \in \mathcal{I}$ and any strategy $s'_I \neq s_I$, the market share m_I satisfies $m_I(s'_I, s_{-I}) \leq$

$m_I(s_I, s_{-I})$. Similarly, $s_{\mathcal{I}}$ is a consumer surplus Nash equilibrium of the game $(M, \mu, s_{\mathcal{N}}, \mathcal{I})$ if for any $I \in \mathcal{I}$ and any strategy $s'_I \neq s_I$, the consumer surplus $\Phi_{\mathcal{I}}$ satisfies $\Phi_{\mathcal{I}}(s'_I, s_{-I}) \leq \Phi_{\mathcal{I}}(s_I, s_{-I})$.

Corollary 1: If $s_{\mathcal{I}}$ is a market shares Nash equilibrium of the oligopolistic game $(M, \mu, \mathcal{N}, \mathcal{I})$, then it is also a consumer surplus $\epsilon_{s_{\mathcal{I}}}$ -Nash equilibrium, where $\epsilon_{s_{\mathcal{I}}} = \max\{\epsilon_{s_I} : I \in \mathcal{I}\}$. Conversely, if $s_{\mathcal{I}}$ is a consumer surplus Nash equilibrium, then it is also a market share $\delta_{s_{\mathcal{I}}}$ -Nash equilibrium, where $\delta_{s_{\mathcal{I}}} = \max\{\delta_{s_I} : I \in \mathcal{I}\}$.

Proof of Corollary 1: If $s_{\mathcal{I}}$ is a market share Nash equilibrium, by definition, each s_I is a market share best response of s_{-I} . Therefore, by Theorem 6, we have

$$\Phi_{\mathcal{I}} \geq \Phi'_{\mathcal{I}} - \epsilon_{s_{-I}} \geq \Phi'_{\mathcal{I}} - \epsilon_{s_{\mathcal{I}}}, \quad \forall s'_I \neq s_I,$$

which concludes that $s_{\mathcal{I}}$ is a consumer surplus $\epsilon_{s_{\mathcal{I}}}$ -Nash equilibrium. If $s_{\mathcal{I}}$ is a consumer surplus Nash equilibrium, by definition, each s_I is a consumer surplus best response of s_{-I} . Therefore, by Theorem 6, we have

$$m_I \geq m'_I - \delta_{s_I} \geq m'_I - \delta_{s_{\mathcal{I}}}, \quad \forall s'_I \neq s_I,$$

which concludes that $s_{\mathcal{I}}$ is a market share $\delta_{s_{\mathcal{I}}}$ -Nash equilibrium. ■

As a direct consequence of Theorem 6, Corollary 1 addresses that the objectives of maximizing market share and maximizing consumer surplus are also closely aligned under Nash equilibria of the oligopolistic game $(M, \mu, \mathcal{N}, \mathcal{I})$.

Regulatory Implications: In the oligopolistic scenario, all ISPs' optimal strategies are closely aligned with the consumer surplus. Even some ISPs use sub-optimal decisions, any remaining ISPs' optimal strategy would still nearly maximize the system consumer surplus (Theorem 6). This alignment with consumer surplus also sustains under Nash equilibria of the multi-ISP competition game (Corollary 1). In conclusion, network neutral regulations are not needed or should not be imposed under a competition market; otherwise, the achieved consumer surplus would be sub-optimal compared to what can be achieved in the efficient Nash equilibria. However, regulations should enforce the ISPs to be transparent in the sense that each ISP's capacity and strategy should be common knowledge to all ISPs. This would help the market converge to an efficient equilibrium more easily.

V. RELATED WORK

Despite of its short history, plenty of work on network neutrality can be found in computer science [11], [20], [6], [12], [23], [17], economics [8], [15], and law [25], [24] literature.

From an economics perspective, Sidak [24] looked at the network neutrality regulation from consumer welfare's point of view and argued that differential pricing is essential to the maximization of welfare. We also focus on the consumer welfare and seek the conditions under which ISPs' strategy would be aligned with consumer welfare. Choi et al. [8] analyzed the effect of neutral regulations on ISPs' investment incentive and found that capacity expansion decreases the sale price of the premium service. This coincides with our finding under the monopolistic scenario; however, under oligopolistic competitions, we find that ISPs do have incentives to increase capacity so as to maximize market share.

From an engineering perspective, Dhamdhere et al. [12] took a profitability perspective and concluded that the ISPs can still be profitable without violating network neutrality. This supports our proposal of a Public Option ISP that can be implemented and sustained by either a government or a private organization. Crowcroft [11] reviewed the technical aspects of network neutrality and concluded that "perfect" network neutrality has never been and should not be engineered. We share the same view that under competition, network neutrality regulation is not necessary; while, under a monopolistic market, a non-regulatory alternative can be a Public Option ISP that incentivizes the existing ISP to maximize consumer surplus.

From a modeling point of view, Musacchio et al. [20] considered advertising CPs and also used a two-stage model under which ISPs move first. Their focus was primarily on a monopolistic ISP. Caron

et al. [6] modeled differentiated pricing for two application type. Our model captures more applications types with parameters of popularity (α_i), maximum throughput (θ_i) and sensitivity to the throughput (β_i). Another departure in our approach from previous analyses is the way we model traffic and congestion in the network. Traditionally, the classical $M/M/1$ formula for delay has been used to abstract out traffic and congestion [8] in economic analyses. Our view is that a more appropriate approach is to more faithfully model closed loop protocols like TCP that carry most of the traffic on the Internet. Shetty et al. [23] used a similar PMP-like two-class service differentiations and considered capacity planning, regulation as well as differentiated pricing to consumers. Our differentiated pricing focuses on the CP-side, where the CPs choose service classes and consumers choose ISPs.

From a regulatory aspect, Wu [25] surveyed the discriminatory practices, e.g. selectively dropping packets, of broadband and cable operators and proposed solutions to manage bandwidth and police ISPs so as to avoid discrimination. Shetty et al. [23] proposed a simple regulatory tool to restrict the percentage of capacity the ISPs dedicate to a premium service class. Economides et al. [13] compared various regulations for quality of service, price discrimination and exclusive contracts, and drew conclusions on desirable regulation regimes. Ma et al. [17] considered the ISP settlement aspect and advocated the use of Shapley value as profit-sharing mechanism to encourage ISPs to maximize social welfare. Our proposal of a Public Option ISP, on the other hand, is a non-regulatory alternative to the network neutral regulations.

VI. DISCUSSION AND CONCLUSIONS

In a monopolistic market, the ISP's selfish non-neutral strategy hurts consumer surplus. Although network neutral regulation can improve consumer surplus under that case, we find a better non-regulatory alternative which is to introduce a Public Option ISP. The existence of a Public Option ISP incentivizes the existing ISP's strategy to be aligned with consumer surplus, and achieve higher consumer surplus than that under network neutral regulations. In an oligopolistic competition, market force influences ISPs' non-neutral strategies to be aligned with consumer surplus and ISPs will get market shares proportional to their capacities. Under this case, the existence of a Public Option ISP would be sub-optimal compared to the efficient Nash equilibria; however, its damage is very limited because the Public Option ISP would be the only one that uses a sub-optimal strategy, where all other ISPs can adapt to optimal strategies and more consumers will move from the Public Option to better and non-neutral ISPs. Of course, there is no reason why the Public Option ISP cannot perform the price discrimination that aligns with the consumer surplus, which induces an efficient Nash equilibrium in theory. However, implementing a neutral public option will avoid mistakes or accidental "collusion" with the existing ISPs in the market. In contrast, if network neutral regulation is enforced, all ISPs will have to perform an neutral but inefficient strategy, which could reduce the consumer surplus substantially.

Theoretically speaking, the existence of a Public Option ISP will be effective if $\mu_{PO} > 0$, regardless how large its capacity is. This is because, in the idealized game model, we assume that an ISP's sole objective is to maximize its market share. In practice, ISPs will trade off its market share with potential revenue from the CPs, which depends on the characteristics of the CPs, e.g. their profit margin and throughput sensitivity, and the condition of the system, e.g. the available capacity and congestion level. Moreover, ISPs might be able to use the CP-side revenue to subsidize the service fees for consumers so as to increase market share.

We envision that the Public Option should be implemented as the safety net, or the last/back-up choice, for the consumers if the existing commercial ISPs' strategy hurt consumer surplus. The more ISPs competing in a market, the less the market needs a public option and the less capacity we need to deploy for the Public Option ISP to be effective. In the most hostile case where only one monopolistic ISP exists in the market, a Public Option ISP could be effective as long as it has a capacity that is larger than the percentage of consumers that the monopoly cannot afford to lose. For example, if 10% of the market share is critical for the monopoly, implementing 10% of its capacity would be able to at least "steal"

10% of consumers from the monopoly if it follows a network neutral strategy. If the monopoly applies a worse than neutral strategy for consumer surplus, it will lose even more. In that sense, although 10% of the capacity will not be operating optimally, its existence incentivizes the remaining 90% maximizing for consumer surplus, which could result in much better consumer surplus than requiring the monopoly to follow network neutral regulations.

Summarizing, we believe our paper sheds new light on and informs the continuing debate on the role of regulation on the Internet and our introduction of the Public Option ISP is an important contribution.

REFERENCES

- [1] FCC Acts to Preserve Internet Freedom and Openness. News Release(12/21/10). http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db1221/DOC-303745A1.pdf.
- [2] Google corporate website. <http://www.google.com/corporate/tech.html>.
- [3] Local Loop Unbundling. http://en.wikipedia.org/wiki/Local-loop_unbundling.
- [4] Netflix technology blog. <http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html>.
- [5] M. Campbell and J. Browning. Apple, Google asked to pay up as mobile operators face data flood. *Bloomberg News*, December 7 2010.
- [6] S. Caron, G. Kesidis, and E. Altman. Application neutrality and a paradox of side payments. *Proceedings of the Re-Architecting the Internet Workshop (ACM ReARCH '10)*, November 2010.
- [7] D. M. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems*, 17(1).
- [8] J. P. Choi and B.-C. Kim. Net neutrality and investment incentives. *The Rand Journal of Economics*, 41(3):446–471, Autumn 2010.
- [9] D. Clark. The design philosophy of the DARPA Internet protocols. *ACM SIGCOMM Computer Communication Review*, 18(4), August 1988.
- [10] C. Courcoubetis and R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. John Wiley & Sons Ltd., 2003.
- [11] J. Crowcroft. Net neutrality: the technical side of the debate: a white paper. *ACM SIGCOMM Computer Communication Review*, 37(1), January 2007.
- [12] A. Dhamdhere and C. Dovrolis. Can ISPs be profitable without violating network neutrality? *Proceedings of 2008 ACM Network Economics (NetEcon)*, August 2008.
- [13] N. Economides and J. Tag. Network neutrality and network management regulation: Quality of service, price discrimination, and exclusive contracts. *Research Handbook on Governance of the Internet*. Ed. Ian Brown. London: Edward Elgar, 2012.
- [14] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr. Complexity of Internet interconnections: Technology, incentives and implications for policy. *The 35th Research Conference on Communication, Information and Internet Policy (TPRC)*, 2007.
- [15] B. Hermalin and M. L. Katz. The economics of product-line restrictions with an application to the network neutrality debate. *Information Economics and Policy*, 19(2):215–248, June 2007.
- [16] C. Labovitz, D. McPherson, S. Iekel-Johnson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *Proceedings of the ACM SigComm*, New Delhi, India, 2010.
- [17] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. Internet Economics: The use of Shapley value for ISP settlement. *IEEE/ACM Transactions on Networking*, 18(3), June 2010.
- [18] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [19] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5), October 2000.
- [20] J. Musacchio, G. Schwartz, and J. Walrand. Network neutrality and provider investment incentives. *Asilomar Conference*, pages 1437–1444, November 2007.
- [21] A. Odlyzko. Paris metro pricing for the Internet. *Proceedings of ACM Conference on Electronic Commerce (EC'99)*, pages 140–147, 1999.
- [22] M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press Cambridge, Massachusetts, 1994.
- [23] N. Shetty, G. Schwartz, and J. Walrand. Internet QoS and regulations. *IEEE/ACM Transactions on Networking*, 18(6), December 2010.
- [24] J. G. Sidak. A consumer-welfare approach to network neutrality regulation of the Internet. *Journal of Competition Law and Economics*, 2(3):349–474, September 2006.
- [25] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 141, 2005.

APPENDIX

We also perform another set of experiments of the same set of 1000 CPs used in our experiments in Section III and IV. The only difference is that the parameter ϕ_i has a distribution $\phi_i = U[0, U[0, 10]]$, which is independent of β_i . Figure 9 and 10 illustrate the consumer surplus under $\kappa = 1$ and various strategies of the ISP respectively. Since the characteristics of the CPs are the same as our previous experiments, the CPs' decision and the ISP's revenues are the same as before.




Fig. 9. Per capita surplus Φ under $\kappa = 1$.




Fig. 10. Φ under various strategies $s_I = (\kappa, c)$ versus per capita capacity ν .

Figure 11 and 12 show the results that are parallel to figure 7 and 8. We still find the same observations as in our previous experiments.




Fig. 11. ISP I 's market share m_I and per capita surplus Ψ_I and per capita consumer surplus Φ under $\kappa = 1$.

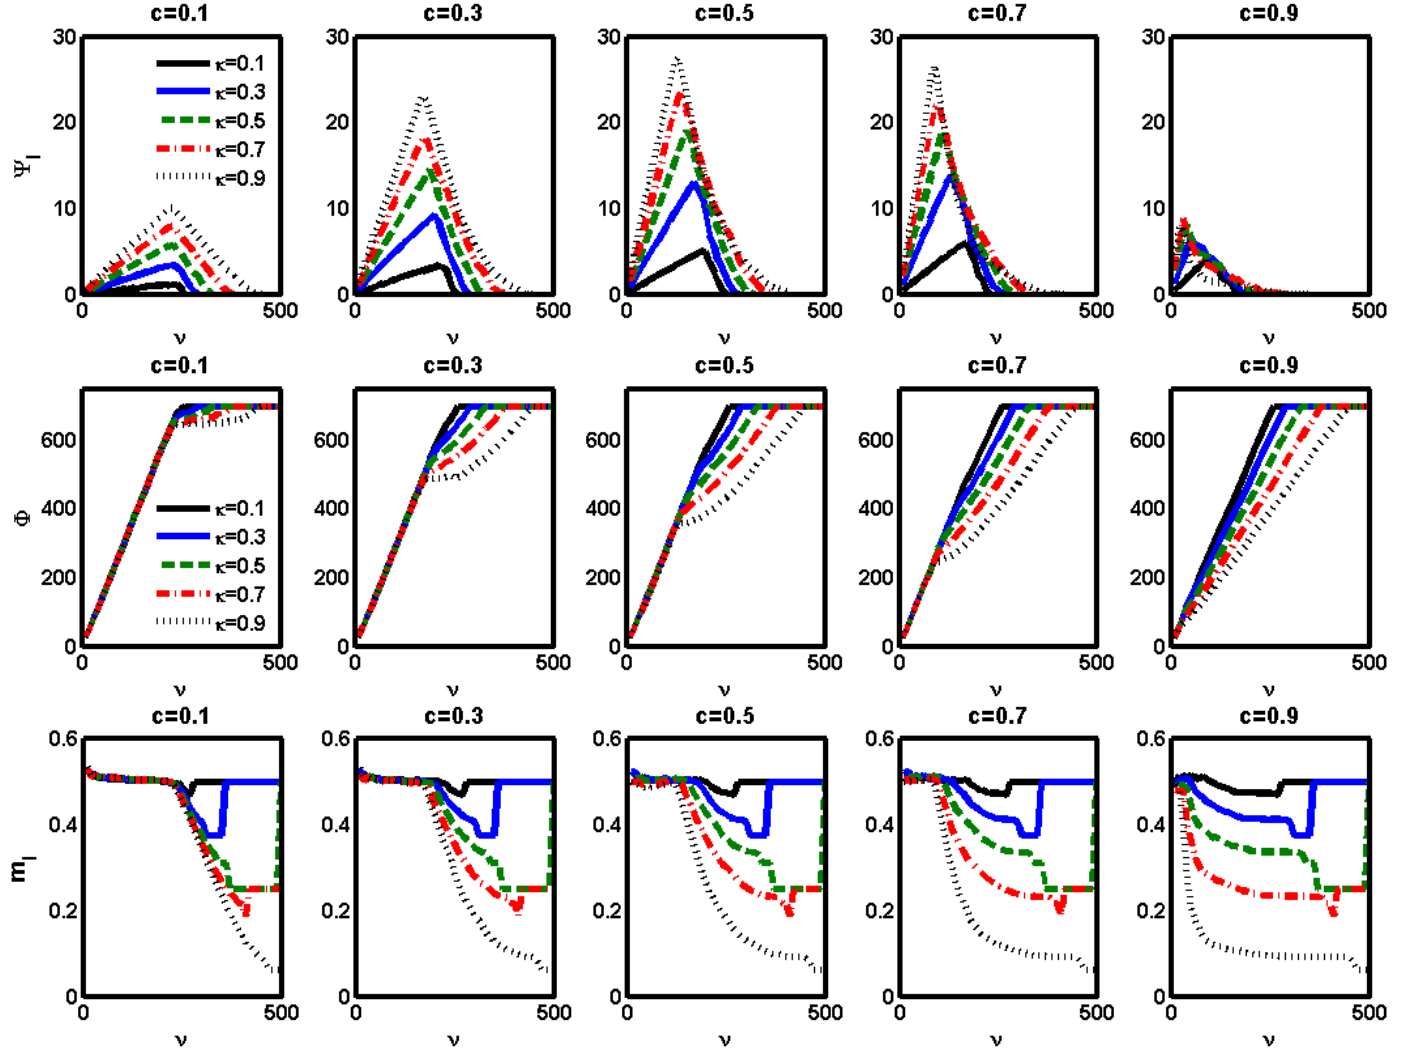


Fig. 12. ISP I 's market share m_I and per capita surplus Ψ_I and per capita consumer surplus Φ under $\kappa = 1$.